

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 7

### REMARKS

The present response is intended to be fully responsive to all points of objection and/or rejection raised by the Examiner and is believed to place the application in condition for allowance. Favorable reconsideration and allowance of the application is respectfully requested.

Applicants assert that the present invention is new, non-obvious and useful. Prompt consideration and allowance of the claims is respectfully requested.

### Objections to the Specification

The specification has been objected to because the word "will" on page 6 line 22 is a typographical error. The error has been corrected (see attached amendments to the specification on page 9). Applicants therefore respectfully request that the objection be withdrawn.

### Status of Claims

Claims 1-14 are pending in the application.

Claims 1-14 have been rejected.

### Claim Rejections

#### *35 U.S.C. § 112 Rejections*

In the Office Action, the Examiner rejected claim 4 under 35 U.S.C. § 112, second paragraph, as being vague and indefinite due to the word "substantially". The Examiner further objected to Claim 4 under 37 CFR 1.75 (c) as being of improper dependent form for failing to further limit the subject matter of previous claim 1. Claim 4 has been amended to remove the word "substantially". Claim 4 has also been amended to recite "said copying" instead of "said

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 8

transferring" because the term transferring is not used in claim 1 on which claim 4 depends. See attached claim listing beginning on page 10. With these amendments, it should now be clear that claim 4 further limits claim 1 in that claim 4 adds the limitation that the copying is performed immediately after the backing up.

Based on these amendments, Applicants respectfully assert that claim 4 is proper under 35 USC 112 and 37 CFR 1.75 (c) and request that the rejection and objection be withdrawn.

### ***35 U.S.C. § 103 Rejections***

In the Office Action, the Examiner rejected claims 1, 4, 6, 13-14 under 35 U.S.C. § 103(a), as being unpatentable over Cooke, Jr. et al. (U.S. Pat No. 6,574,629, hereinafter "Cooke") in view of Otterness et al (U.S. Pat. No. 6,460, 122, hereinafter "Otterness").

Applicants respectfully traverse the rejection of 1, 4, 6, 13-14 under 35 U.S.C. § 103(a), as being unpatentable over Cooke in view of Otterness.

Applicants respectfully assert that neither Cooke nor Otterness, singly or in combination, teaches or suggests a storage system which includes both a high performance high reliability storage medium and a high performance low reliability storage medium.

The current invention discloses a storage system including inter-alia a high performance high reliability storage medium and a high performance low reliability storage medium and a method for managing the system:

"In accordance with the current invention, an inexpensive storage system is disclosed along with methods of managing such a system.. In one preferred embodiment, the system includes a high performance high reliability storage medium configured for initial storage of data, a low performance high reliability storage medium configured for backup of data initially stored on the high performance high reliability storage medium, and a high performance low reliability storage medium, configured to receive data transferred from the high performance high reliability storage medium, after the data has been backed up on the low performance high reliability storage medium. This significantly reduces the cost of the system without substantially

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 9

comprising performance. Reliability is likewise maintained owing to the high reliability backup." (Abstract)

By way of contrast, Cooke discloses in column 9, lines 9 to 40, a magneto optical disk MOD based Archive Station:

"The MOD-based archive station includes a three-tiered storage system that uses hard disks (preferably 4 giga-bytes each) for on-line data storage (usually one week's worth of data), MODs (i.e., high-capacity, removable storage media having preferable either 2.3 or 4.6 giga-bytes of storage each) for intermediate-term data storage (usually one to two years), and DLT media for long-term backup storage. Additional RAID ("Redundant Array of Inexpensive Disks") may also be provided for expanded short term storage, depending upon the storage requirements of the facility and the number of studies handled thereby. The portion of short term storage, i.e., the RAID and other hard disks, which stores images comprises the MOD-based archive's cache. ... In preferred embodiments of the invention, the workstation also runs an "auto-pilot" routine which patrols the archive's cache and which automatically deletes studies once they have been archived on a MOD. This is done in order to prevent the cache from "overflowing". "

Cooke also discloses in Column 9, lines 44-64, an alternative archive station, the digital linear tape DLT based Archive Station:

"In the MOD-based archive station described above, DLT is configured to backup the primary long-term storage in order to provide data redundancy. On the other hand, in the DLT-based archive station, the DLT is used for primary intermediate and long-term storage. In preferred embodiments, DLTs having a storage capacity of 35 giga-bytes each are used; although the invention can be used with DLTs having other storage capacities as well.... In addition, MODs may also be provided for additional and/or backup intermediate and long-term storage. As above, RAID and hard disks are also provided for short-term storage of studies and for database storage. The DLT-based archive station provides both manual and automatic control of data migration from short term storage to DLT. These controls are substantially similar to those described above with respect to the MOD-based archive station and, therefore, are not repeated here for the sake of brevity."

Applicants assert that RAID (used in Cooke for short term storage in both the MOD based archive station and the DLT based archive station) is known in the art to be a high reliability high performance storage medium. For example, in the background

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 10

section of the current application, on page 2, lines 1-8, RAID is defined as a high reliability high performance storage medium:

"For example, a system which uses a Redundant Array of Inexpensive Disks also known as Redundant Array of Independent Disks (RAID) stores the same data in different places (thus, redundantly) on multiple hard disks. By placing data on multiple disks, I/O operations can overlap in a balanced way, improving performance. Since the use of multiple disks increases the mean time between failure, storing data redundantly also increases fault-tolerance. Therefore usage of a **RAID provides relatively high reliability via redundancy, while preserving high performance**, but also increasing cost." (Emphasis added).

Applicants further assert that both the MOD based archive station of Cooke and the DLT-based archive station of Cooke resembles the prior art heterogeneous storage system discussed with reference to Fig 1 in the current application.

For example, a typical prior art heterogeneous storage system includes expensive, high performance, high reliability magnetic disk storage (for example RAID) plus low cost, low performance, high reliability secondary storage (for example tapes or magnetic/optics such as Magneto Optical (MO), Compact Disc (CD), Digital Versatile Disc also known as Digital Video Disc (DVD) etc.)

Fig. 1 depicts a prior art heterogeneous storage system 100 which allows storage and retrieval of data by client(s) 110.... Client 110 stores data in heterogeneous storage system 100... The data is initially stored in a high-reliability high-performance storage 120. It is assumed that the data is eventually backed up on secondary (high-reliability low-performance) storage 140... Once the data has been backed up, it is possible that the data will be eliminated immediately from high performance storage 120, or will later be eliminated from high-performance storage 120, for example once storage 120 is full and new data is sent for storage. In effect, the secondary storage 140 archives the data when it is no longer likely to be required on an ongoing basis and, to this extent, is off-line and therefore not directly accessible. (Page 2, line 11 to page 3, line 2).

Both the archive stations described in Cooke (MOD and DLT) and the prior art heterogeneous storage system described in the current application include high performance high reliability storage, such as RAID. In addition, both the archive stations described in Cooke

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 11

(MOD and DLT) and the prior art heterogeneous storage system described in the current application include low performance high reliability secondary storage such as magneto optics and/or tapes. Moreover, neither the archive stations described in Cooke (MOD and DLT) nor the prior art heterogeneous storage system described in the current application give even a hint or a suggestion of the usage of low reliability storage.

Cooke also mentions the usage of other hard disks for (high performance) short term storage in column 9 lines 9 to 40, and column 9, lines 44-64 reproduced above. There is no mention in Cooke of any practical differentiation between the hard disks and RAID within the MOD based archive station. Both the hard disks and RAID are referred to as the "MOD-based archive's cache" in column 9, line 21. Also, both the hard disks and RAID are subject to an auto-pilot routine which automatically deletes studies once they have been archived on a MOD (see column 9, lines 36-38). In addition both the hard disks and RAID are subject to the same manual and automatic control of data migration from short term storage to DLT (see column 9, lines 57-60). Moreover, there is no mention in Cooke of a desirability for transfer between the hard disks and RAID, reinforcing the interchangeability of the two within the archive station. Certainly, there is neither a hint nor a suggestion in Cooke of a differentiation between the hard disks and RAID in terms of reliability. It may therefore be assumed that just as the RAID used in Cooke for short term storage is a high performance high reliability storage medium, the other hard disks used in Cooke for short term storage must also comprise high performance high reliability storage media. In any case, what is certainly clear and not subject to speculation is that there is no suggestion that the other hard disks used in Cooke for short term storage are anything other than high performance high reliability storage media.

Otterness et al teaches a multi-level cache architecture to be used in either a multi processor or in a multi controller storage environment (column 3 lines 61-63). The caches are differentiated by their performance characteristics (time, speed).

"For the purposes of this description, data which is immediately accessible to the processor, i.e., data which does not need to be copied from another location, read

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 12

from disk, or the like, will be referred to as level 0 (primary) cache. Data which is stored on a different controller, in a remote cache, or on a different medium, but can be accessed by the primary processor through copy or move operations will be referred to as level 1 to level n cache. All caches besides the level 0 cache will be collectively referred to as the secondary cache. For the purpose of this discussion there can be multiple pools of level 1 to level n cache **all which may require different amounts of time** to move data from the secondary cache to the level 0 cache." (Column 4 line 42-54, emphasis added)

"For the purposes of this description, the various caches are referred to as various levels from the point of view of (or relative to) a particular controller. The "level 0" cache is considered the first level of interaction with host operations. Data is moved to and from the level 0 cache to the host computer system. The level 0 cache is also the **fastest** temporary storage medium available to the storage controller. The level 1 cache is a **second fastest** temporary storage medium, it could be an alternate controller, or a solid-state disk, or the like. Further cache levels, for example, level 2, level 3, . . . , level n, indicate that **longer access time and a slower path** to reach the host system with the requested or received data." (Column 20, lines 55-67, emphasis added)

As seen from the paragraphs quoted above, the differentiation between the caches in Otterness is based on performance. There is neither a hint nor a suggestion in Otterness of a differentiation among caches based on other factors that reflect or impact on reliability. Specifically, there is neither a hint nor a suggestion in Otterness of a storage system which includes both a high performance high reliability storage medium and a high performance low reliability storage medium.

Applicants respectfully assert that based on the arguments presented above, neither Cooke nor Otterness singly or in combination teaches or suggests a storage system which includes both a high performance high reliability storage medium and a high performance low reliability storage medium. All the more so, neither Cooke nor Otterness singly or in combination teaches or suggests independent claims 1, 6, 13, and 14. All four independent claims recite the usage of a high performance high reliability storage medium and a high performance low reliability storage medium, where the high performance low reliability storage medium receives data from the high performance high reliability storage medium which has been previously backed up on a low performance high reliability storage medium.

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 13

In claim 1:

"...after said backing up, copying at least some of said data from said high reliability high performance storage medium to a low reliability high performance storage medium..."

In claim 6:

"... a high performance low reliability storage medium configured to receive data transferred from said high performance high reliability storage medium, after said data has been backed upon said low performance high reliability storage medium."

In claim 13:

"...after said backing up, copying at least some of said data from said high reliability high performance storage medium to a low reliability high performance storage medium..."

In claim 14:

"...computer readable program code for causing the computer after said backing up to copy at least some of said data from said high reliability high performance storage medium to a low reliability high performance storage medium..."


Based on the above, Applicants respectfully assert that independent claims 1, 6, 13, and 14 are allowable. Claims 2, 3, 4 (as amended), 5, 7-12 depend from, directly or indirectly, claims 1, 6, 13, and 14 and therefore include all the elements of those claims. Therefore, Applicants respectfully assert that claims 2-4, 5, 7-12 are likewise allowable. Accordingly, Applicants respectfully request that the Examiner withdraw the rejections to independent claims 1, 6, 13 and 14, and to claims 2-4, 5, 7-12 dependent thereon.

Applicants believe the remarks presented hereinabove to be fully responsive to all of the grounds of rejection raised by the Examiner. In view of these remarks, Applicants respectfully submit that the specification and all of the claims in the present application are in order for allowance. Notice to this effect is hereby requested.

APPLICANTS: Melament et al  
SERIAL NO.: 10/697,186  
FILED: October 30, 2003  
Page 14

Please charge any fees associated with this paper to Deposit Account No. 09-0468.

Respectfully submitted,

By:   
Stephen C. Kaufman  
Reg. No. 29,551  
Phone No. (914) 945-3197

Date: March 13, 2006

IBM Corporation  
Intellectual Property Law Dept.  
P. O. Box 218  
Yorktown Heights, New York 10598